

ing energy and memory budgets. Inasmuch as observation opportunities repeat, the theoretical framework for evaluation of candidate solutions includes a cycle bound.

The method includes the use of an iterative optimization algorithm known in the art the “squeaky wheel optimization.” This algorithm consists of the following steps:

1. Sort all of the target points according to priority.
2. Schedule each target point, greedily, in order of priority.
3. Increase priorities for those points that were omitted from the schedule in step 2.
4. Iterate by repeating from step 1.

Thus, points that are not initially scheduled have a greater probability of being scheduled on subsequent iterations because each time a point is not scheduled,

its priority is increased, until eventually it becomes the first point scheduled.

At each iteration, points are considered in order of priority and observations considered in descending order of the contribution of each to coverage until a minimum acceptable level of coverage of each affected point has been scheduled. Only observations that cannot oversubscribe memory or energy and that respect transition duration constraints are considered. The asymptotic time complexity (in effect, the time needed to perform the computations) is of the order of a number proportional to $n \log(n)$ per iteration, where n is the number of points. This computation time is basically proportional to the time needed to sort the points according to priority.

For reasoning about multiple cycles, multiple cycles are used in representing

energy and memory, but a single cycle is used in representing observations. The single-cycle representation can be characterized as compressed in that, relative to a time-line representation, it requires less memory to represent all possible observations over all cycles. Instead of listing all observations individually, one lists a single cycle of observations with labels that represent which observations belong to which cycles. The amount of memory needed to encode observations in this approach is proportional to the total number of observations.

This work was done by Steve Chien and Russell Knight of Caltech for NASA’s Jet Propulsion Laboratory.

The software used in this innovation is available for commercial licensing. Please contact Karina Edmonds of the California Institute of Technology at (626) 395-2322. Refer to NPO-43768.

Self-Supervised Learning of Terrain Traversability From Proprioceptive Sensors

This system enables a vehicle to scan its surroundings and adapt to conditions by learning about them on the fly.

NASA’s Jet Propulsion Laboratory, Pasadena, California

Robust and reliable autonomous navigation in unstructured, off-road terrain is a critical element in making unmanned ground vehicles a reality. Existing approaches tend to rely on evaluating the traversability of terrain based on fixed parameters obtained via testing in specific environments. This results in a system that handles the terrain well that it trained in, but is unable to process terrain outside its test parameters.

An adaptive system does not take the place of training, but supplements it. Whereas training imprints certain environments, an adaptive system would imprint terrain elements and the interac-

tions amongst them, and allow the vehicle to build a map of local elements using proprioceptive sensors. Such sensors can include velocity, wheel slippage, bumper hits, and accelerometers. Data obtained by the sensors can be compared to observations from ranging sensors such as cameras and LADAR (laser detection and ranging) in order to adapt to any kind of terrain. In this way, it could sample its surroundings not only to create a map of clear space, but also of what kind of space it is and its composition.

By having a set of building blocks consisting of terrain features, a vehicle can adapt to terrain that it has never seen be-

fore, and thus be robust to a changing environment. New observations could be added to its library, enabling it to infer terrain types that it wasn’t trained on. This would be very useful in alien environments, where many of the physical features are known, but some are not. For example, a seemingly flat, hard plain could actually be soft sand, and the vehicle would sense the sand and avoid it automatically.

This work was done by Max Bajracharya, Andrew B. Howard, and Larry H. Matthies of Caltech for NASA’s Jet Propulsion Laboratory. For more information, contact iaoffice@jpl.nasa.gov. NPO-46601